

Enhancing Fault Tolerance Capabilities in Quorum-based Cycle Routing

The final publication is available at IEEE via <http://dx.doi.org/10.1109/RNDM.2015.7324305>

Cory J. Kleinheksel, Member, IEEE and Arun K. Somani, Fellow, IEEE
Electrical and Computer Engineering, Iowa State University, Ames, Iowa 50011
Telephone: (515) 294-0442, Fax: (515) 294-9273, Email: {cklein, arun}@iastate.edu

Abstract—In this paper we propose a generalized R redundancy cycle technique that provides optical networks almost fault-tolerant communications. More importantly, when applied using only single cycles rather than the standard paired cycles, the generalized R redundancy technique is shown to almost halve the necessary light-trail resources in the network while maintaining the fault-tolerance and dependability expected from cycle-based routing.

For efficiency and distributed control, it is common in distributed systems and algorithms to group nodes into intersecting sets referred to as quorum sets. Optimal communication quorum sets forming optical cycles based on light-trails have been shown to flexibly and efficiently route both point-to-point and multipoint-to-multipoint traffic requests. Commonly cycle routing techniques will use pairs of cycles to achieve both routing and fault-tolerance, which uses substantial resources and creates the potential for underutilization. Instead, we intentionally utilize R redundancy within the quorum cycles for fault-tolerance such that every point-to-point communication pairs occur in at least R cycles. The result is a set of $R = 3$ redundant cycles with 93.23 - 99.34% fault coverage even with two simultaneous faults all while using 38.85 - 42.39% fewer resources.

I. INTRODUCTION

We developed a novel method to deliver the almost fault-tolerant capabilities of cycles in an optical network and the potential for significantly reducing the resource utilization when compared to the state-of-art techniques. Cycle-based routing can satisfy both dynamic point-to-point and multi-point optical communications. Cycles are created using quorums of nodes. Within a cycle, multicasts to all nodes in that cycle is possible. The quorum intersection property and the use of cyclic quorums sets provide all of the unicast capabilities. Exploiting the same properties we can achieve efficient broadcasts with $O(\sqrt{N})$ multicasts.

Optical networks are depended upon for high speed communications in distributed algorithms, as much as they are needed for the arbitrary point-to-point communications. Failures within a network are to be expected and can happen as much as every couple days. Protecting against these optical circuit faults is critical and there are many different approaches depending on the network needs and individual circumstances.

For efficiency and distributed control, it is common in distributed systems and algorithms to group nodes into intersecting sets referred to as quorum sets. Quorums sets for cycle-based routing to efficiently support arbitrary point-to-point and multi-point optical communication were first proposed in [1] with fault-tolerance analyzed in [2]. In this paper we apply

the established quorum set theory and generalize additional requirements proposed in [3] to form suitable R redundant quorums for our optical network routing.

The rest of the paper is organized as follows. Sections II, III, and IV establish the network model, node communication, and path routing / fault-tolerance. In Section V, we discuss our application of the distributed efficiency of the quorum sets to routing optical cycles. We also define and generalize R redundant quorums sets for this application. Section VI analyzes the performance and fault-tolerance of our redundant quorums set cycle routing techniques for the state of art paired cycle routing solutions. Lastly, Section VII analyzes our proposed single cycle routing solutions and compares them to the more resource intensive paired cycle routing solution.

II. NETWORK MODEL

No two fiber-optic networks are the same. Some stretch hundreds of kilometers, while other networks are contained within buildings or rooms. Regardless of the physical environment, these optical circuits are depended upon for high speed communications. Thus it is important to extract the network's critical components that affect its ability to deliver reliable, arbitrary point-to-point and multi-point communications.

These fiber-optic networks consist of several transmitters and receivers interconnected by fiber-optic cables. As you might expect, transmitters and receivers are typically found together and can generically be called an optical node. The cables form the links (i.e. edges) between those nodes, which leads to a convenient model of a network in terms of a graph $G = (V, E)$. V is the set of nodes in the network and E is the set of edges.

Edge (a_i, a_j) is a fiber-optic link connecting nodes a_i and a_j in the network, where $a_i, a_j \in V$ and $(a_i, a_j) \in E$. It is a general assumption that the same set of optical wavelengths are available on all edges in E . The number of wavelengths available per optical fiber is dependent on the fiber-optic cables and the transmitter/receiver pairs.

III. LIGHT-TRAILS

Lightpaths were a critical building block in the first optical communications, but required significant traffic engineering and were not able to support multicast traffic. Light-trails were proposed in [4] as a solution to these challenges and could be built using commercial off-the-shelf technology.

Light-trails enable fast, dynamic creation of an unidirectional optical communication channel. This communication




Fig. 1. Four nodes in a light-trail architecture.

channel allows for channel receive and transmit access to all connected nodes, making them more suitable for IP centric traffic [5]. Point-to-point communications from an upstream node to a downstream node can be scheduled on the shared light-trail. Similarly, an upstream node can multicast to any number of downstream nodes.

An example four-node light-trail can be seen in Fig. 1. Optical shutters isolate an optical signal to a specific light-trail and allow for wavelength reuse within the network. Start and end nodes have their optical shutters in the *off* state, while intermediate nodes have their optical shutters in the *on* state. The communication is all optical from the start node to the end node. Being all optical avoids any energy inefficiencies and time delays associated with unnecessary Optical-to-Electrical-to-Optical (O/E/O) conversions at intermediate hops.

Nodes can receive from the incoming signal while the signal is simultaneously continuing to downstream nodes, sometimes referred to as a drop and continue function. Early technology supported only a few wavelengths, however the latest devices may support over 100 channels, hence allowing multiple light-trails to share the same edge in the network for a combined over 1-Terabits/s.

A scheduling protocol is in place to avoid collisions within a light-trail and controls when nodes are able to transmit to downstream nodes. The scheduling is generally assumed to occur over a control channel, which may or may not be separate from the shared optical fiber that is being used for the light-trail.

IV. LIGHT-TRAIL, CYCLE ROUTING, AND FAULT-TOLERANCE

Point-to-point and multi-point traffic requests have a set of nodes $C = \{a_i, \dots, a_j\}$ that wish to communicate. Failures within an optical network are to be expected, hence requests need to be protected against network faults. Establishing a primary and backup multicast path from every node to every other node in C can be a waste of resources though.

The generalization of p-cycle protection to allow for path and link protection was proposed by [6]. P-cycle protection of unicast and multicast traffic networks requires preconfiguration and the offline nature allows for the efficient cycles to be selected [7], [8]. The use of path-pair protection, linked-based shared protection, spanning paths, and p-cycles to protect multicast sessions have all been proposed for WDM networks as well [9]–[13].

The Optimized Collapsed Rings (OCR) single link protection heuristic was developed to address the heterogeneous, part multicast / part unicast, nature of WDM traffic [14]. The Multi-point Cycle Routing Algorithm (MCRA) uses bidirectional cycles for fault-tolerance and is capable of supporting SONET




Fig. 2. Cycle formed using the light-trail architecture.

rings and p-cycles [15]. Although finding the smallest cycle supporting the multi-point communication is NP-Complete, the authors were able to show that their heuristic performed within 1.2 times of the optimal cycle size. ECBRA is a significant improvement of MCRA and outperforms the OCR heuristic [16].

The ECBRA heuristic balances optimality and speed, taking $O(|E| |C|^3)$ steps to find a close to optimal cycle. First, a modified breadth first search is performed on each node in set C of required communication nodes. The goal is to find a shortest path in G that also has the best ratio of nodes from set C vs. total nodes on the path. To complete cycle using the chosen shortest path, a path from the sink node returning to the source node must be found. No links may be used twice. If all nodes in C are in the cycle, then the cycle search is complete. If needed, the final step iteratively removes edges from the cycle and inserts paths through missing nodes in C . Because insertion of the node can be cheaper by removing certain links from the cycle rather than others, the optimal edge removal from the cycle and path replacement is computed for each missing node insertion.

In this work, we utilize the light-trail architecture in the form of a cycle. Figure 2 is simply a light-trail where the start and end node is the same node, referred to as the *hub node* of the cycle. The hub node has its optical shutters in the *off* state, while intermediate nodes have their optical shutters in the *on* state. The resources at each hub node can be utilized to allow all-to-all communication on the cycle using only one light-trail. Traffic from a node to nodes downstream requires a single transmission. Traffic from a node to an upstream node must undergo Optical-to-Electrical-to-Optical (O/E/O) conversion at the hub node and be transmitted on the light-trail a second time.

Alternately, the state of art solutions setup two light-trails, one in each direction. This enables upstream communications without the energy inefficiencies and time delays associated with O/E/O conversions.

V. QUORUMS

In distributed communication and algorithms, coordination, mutual exclusion, and consensus implementations have grouped N nodes into small sets called quorums. This organization of nodes can minimize communications in operations like negotiating access to a global resource.

A quorums set minimally has the property that all quorums in the set must intersect. For distributed implementations, it can also be desirable that each node have equal work and equal responsibility [17].

Not every grouping of nodes into sets (quorums) will result in having these three properties, nor will the quorum sizes be

minimal. [17] proved the lower bound on the size of quorums set having these three properties. Cyclic quorums sets with these properties, like other quorums sets in general, are difficult to find and require an exhaustive search [17], [18]. Quorums sets were shown to efficiently support arbitrary point-to-point and multi-point communication for cycle-based routing in [1], [2]. Redundant quorum sets significantly reduced the required amount of network resources while still maintaining a similar level of fault-tolerance [3].

A. Defining Quorums Set

A is a set of $N = |V|$ nodes. A set S_i is a subset of A . When set Q of subsets (Eq. 2) covers all nodes in A (Eq. 3) and all subsets also have non-empty intersections (Eq. 4), then set Q is called a quorums set.

$$A = \{a_1, \dots, a_N\} \quad (1)$$

$$Q = \{S_1, \dots, S_N\} \quad (2)$$

$$\bigcup_{i=1}^N S_i = \{a_1, \dots, a_N\} = A \quad (3)$$

$$S_i \cap S_j \neq \emptyset, \forall i, j \in 1, 2, \dots, N \quad (4)$$

The lower bounds for the maximum individual quorum size (i.e. $|S_i|$) in a minimum set is K , where Equation 5 holds and $(K - 1)$ is a power of a prime, proved through equivalence to finding a finite projective plane [17]. Additionally it is desirable that each quorum S_i in the quorum set be of equal size (Eq. 6), such that there is equal work and it is desirable that each node be contained in the same number of quorums (Eq. 7), such that there is equal responsibility.

$$N \leq K(K - 1) + 1 \quad (5)$$

$$|S_i| = K, \forall i \in 1, 2, \dots, N \quad (6)$$

$$a_i \text{ is contained in } K S_j's, \forall i \in 1, 2, \dots, N \quad (7)$$

Cyclic quorums adhere to these properties [18]. Cyclic quorums are unique in that once the first quorum (Eq. 8) is defined the remaining quorums can be generated via incrementing (modulus not shown in Eq. 9 for conciseness.) For simplicity assume $a_1 \in S_1$ without loss of generality (any one-to-one re-mapping of entity ids can result in this assumption.)

$$S_1 = \{a_1, \dots, a_j\} \quad (8)$$

$$S_i = \{a_{1+(i-1)}, \dots, a_{j+(i-1)}\} \quad (9)$$

B. Redundant Cyclic Quorums Sets

In this section we define and generalize R redundant quorums sets. The quorum-based cycle solution uses quorums to form a set of communication cycles which were shown to support almost fault-tolerance to fiber optic networks [1], [2]. As defined in Section V-A, there are N quorums in a set and each quorum has K nodes. In analysis of networking capabilities we are interested in whether every node can communicate with every other node. Equation 10 considers the number of communication pairs within a quorum, i.e. the pairs made between K nodes communicating with $(K - 1)$ other nodes in a single quorum. Equation 11 is the total pairs in the size N quorums set solution. For convenience we set

M to be the total pairs for a given network with N nodes and K optimal quorum size.

$$\frac{K(K - 1)}{2} = O(K^2) \quad (10)$$

$$N \frac{K(K - 1)}{2} = O(NK^2) \quad (11)$$

$$M = O(NK^2) \quad (12)$$

When the quorum size, K , is minimal or larger, every pair of nodes (a_i, a_j) occurs together within a quorum in the set at least once. Optical networking however requires all directional point-to-point pairs to exist, i.e. both pairs (a_i, a_j) and (a_j, a_i) . Previously this had been addressed by pairing each cycle with the same cycle with its direction reversed. The key change in [3] is an added requirement that every pair (a_i, a_j) would occur together within at least two quorums rather than just one, which sought to eliminate the need for paired cycles. The number of quorums in the solution remained the same N , hence to create the additional pairs the quorum size had to be enlarged to \hat{K} . Equation 13 calculates the number of node pairs in quorums of size \hat{K} . Equation 14 is our requirement that the total number of pairs have doubled from the original total pairs, M . Finally, Eq. 15 solves for size \hat{K} in relation to optimal K .

$$N \frac{\hat{K}(\hat{K} - 1)}{2} = O(N\hat{K}^2) \quad (13)$$

$$O(N\hat{K}^2) = 2M \quad (14)$$

$$\hat{K} \approx \sqrt{2}K \quad (15)$$

This result is powerful because with a $\sqrt{2}$ factor increase in K the need for paired cycles is reduced and opens the door for considerable resource savings. Still there are applications that may benefit from improved fault tolerance, hence we further generalize this approach for a generic desired R redundant factor to offer an opportunity to enhance the fault tolerance of our quorum-based cycle routing solution. Equation 16 balances the enlarged quorum size \hat{K} solution against the known R times the optimal solution. Equation 17 solves for \hat{K} in terms of known K .

$$O(N\hat{K}^2) = RM \quad (16)$$

$$\hat{K} \approx \sqrt{RK} \quad (17)$$

To the best of our knowledge, no efficient algorithm is known to find quorums of minimum size, particularly with the additional requirement that entity pairs appear a minimum R times within the quorums set solution. [18] used a brute force search to find optimal cyclic quorums for $N = 4 \dots 111$. Using our generalized result from Eq. 17, we too used a brute force search beginning with the smallest possible quorum size for a given number of nodes N and a given desired redundancy factor R .

The resulting redundant quorums were utilized in the following sections, as we analyzed and enhanced the efficacy of quorum-based cycle routing in optical networking.

TABLE I. MEAN LINKS USED BY THE PAIRED CYCLES USING OUR REDUNDANT QUORUM SOLUTION (95% CONFIDENCE INTERVALS.)

Networks	$R = 1$	$R = 2$	$R = 3$
NSFNET	248.96 - 249.88	270.21 - 270.93	289.95 - 290.85
ARPANET	510.03 - 511.24	538.89 - 539.87	587.72 - 588.89
American	641.98 - 643.35	718.87 - 720.19	752.74 - 753.83
Chinese	2673.74 - 2678.28	3053.24 - 3057.70	3270.65 - 3274.51

VI. PAIRED CYCLES NETWORK ANALYSIS

We begin by examining our proposed expansion of redundant quorums by comparing apples-to-apples using the paired cycle routing in prior art. We used four common networks (Fig. 3) and an implementation of the ECBRA heuristic [16] to perform the cycle routing. ECBRA is sensitive to node and edge numbering that a total of 1000 variations on the inputs were considered, each being a one-to-one mapping with the respective network. For simulation of prior art in [1]–[3], we used the $N=4, \dots, 111$ optimal cyclic quorums from [18]. Redundant cyclic quorums for $R = 2$ and $R = 3$ were found using the techniques described in Section V-B.

A. Fault-free operational analysis

It is expected that a majority of the time the optical network will be operating without faults. It is important that the resource utilization during this period be analyzed.

The metric we use to measure resource utilization is the number of links used in a solution. Comparing network-to-network is not particularly insightful, but comparing multiple solutions for a particular network is. The more links that a set of quorum cycles use the fewer (wavelength) resources that can be assigned to each link. Additionally each logical link represents a required physical transmitter and receiver, hence capital costs.

Table I shows that applying redundancy within the quorums and using paired cycles will lead to an increase in mean network links used (95% confidence interval.) $R = 1$, column two, is the standard, no redundant pairs, implementation seen in [1], [2]. $R = 2$ and $R = 3$ have twice and three times redundant pairs present in their quorum solutions respectively. Despite having that added redundancy, the resource usage, columns 3 and 4, only increased 5.63-14.18% and 15.21-22.29% respectively across the networks.

This resource usage result shows that applying our redundant quorums set technique to paired cycle solutions available today will not pose too significant of a resource burden. Next we consider the fault case for paired cycles using our redundant quorum cycle solution and show the increase in resources is being utilized to improve fault recoveries without any optical cycle reconfiguration.

B. Fault-tolerant operational analysis

Optical networks are highly depended upon. The fault-tolerance aspect of this route design is critical. Maintaining the ability to serve all dynamic point-to-point traffic requests despite fault is important.

We assume fiber link failure(s). While a simple model, it does cover most real fault scenarios. This occurs when a link

TABLE II. PERCENT MEAN FAULT COVERAGE (95% CI) OF PAIRED CYCLES USING OUR REDUNDANT QUORUM SOLUTION EXPERIENCING A SINGLE LINK FAULT.

Networks	$R = 1$	$R = 2$	$R = 3$
NSFNET	99.47 - 99.50%	99.83 - 99.84%	99.99 - 99.99%
ARPANET	99.80 - 99.80%	99.93 - 99.94%	99.97 - 99.98%
American	99.62 - 99.63%	99.91 - 99.92%	99.95 - 99.96%
Chinese	99.90 - 99.90%	99.98 - 99.98%	99.99 - 99.99%

is broken, like planned maintenance or the accidental severing during land excavation. Each modeled node has transmitters and receivers that can fail too. Short of a natural disaster, devices will likely fail independently of one another. When a transmitter/receiver pair fails within a modeled node, the effect on the global network is similar to that link failing. Modeling as a single edge failure, while not an exact fault mapping, is an appropriate abstraction.

1) *Single fault case:* To model the fault, we simulate the failure of each edge, $e \in E$, in the network model, $G = (V, E)$. We then examine the network's ability to serve all potential point-to-point requests by counting pairs of nodes that would be able to communicate and conversely those pairs that are unable to communicate. The results are then reported as fault coverage, i.e. total pairs able to communicate as a percentage of total point-to-point pairs. 100% would be perfect coverage, whereas 0% would be no fault coverage at all.

Our simulation results showed our redundant quorum-based cycle technique had 99.83 - 99.98% and 99.95 - 99.99% fault coverage, $R = 2$ and $R = 3$ respectively, in the four networks tested. In Table II, we compare the state-of-art paired cycle approach with our redundant technique also with paired cycles. With single edge failures, the paired cycles had a mean missing communication pair rate of less than 3 pairs or less than 0.53% across all networks (95% CI.) Hence column two, it can be seen that the fault coverage is greater than 99.47%. Our redundant quorum cycles technique, columns 3 and 4, had a mean missing pair rate (95% CI) of less than 0.48 and 0.26 respectively across all networks, which is reflected in fault coverages greater than 99.83%.

Depending on the network, the difference between a 99.5% and 99.99% fault coverage could be significant. Being able to achieve that with only the moderate overheads examined in the previous section is just one of the benefits of the redundant quorums set technique. Being able to dial in on the fault coverage desired using single ($R = 1$), double ($R = 2$), or triple ($R = 3$) redundancy also adds to the flexibility.

2) *Two fault case:* A significantly more complex model considers two faults simultaneously. All possible two edge failure combinations in the network are simulated. We then examine the network's ability to serve all potential point-to-point requests by counting pairs of nodes that would be able to communicate and conversely those pairs that are unable to communicate.

Our simulation results showed our redundant quorum-based cycle technique had 98.65 - 99.91% and 99.04 - 99.95% fault coverage, $R = 2$ and $R = 3$ respectively, in the four networks tested. In Figure 4, we compare the state-of-art paired cycle approach with our redundant technique also with paired cycles. With two edge failures, the paired cycles had a mean missing




Fig. 3. Networks used for simulations: (a) NSFNET, 14-Node/22-Link, (b) ARPANET, 20-Node/31-Link, (c) American Backbone [19], 24-Node/43-Link, and (d) Chinese Backbone [19], 54-Node/103-Link.




Fig. 4. Percent mean fault coverage (95% CI) of paired cycles using our redundant quorum solution experiencing two simultaneous link faults. Confidence intervals are shown for each data point, however tight intervals may be hidden by the data point marker.

communication pair rate of 9 pairs (2.37%) or less across all networks (95% CI.) Hence in the $R = 1$ column, it can be seen that the fault coverage is 97.63% or more. Our redundant quorums set technique, $R = 2$ and $R = 3$, had a missing pair rate (95% CI) of 3.14 and 2.09 or less respectively across all networks, which is reflected in fault coverages of 98.65% or more.

An interesting aspect of the data is that the benefit of additional quorum redundancy in paired cycle solutions is network dependent. All networks had improved fault coverage with additional quorums set redundancy (Fig. 4), however the NSFNET network had larger increases than others. Additionally, all networks increasing to $R = 3$ had diminishing returns compared to the increases seen when moving from $R = 1$ to $R = 2$.

VII. SINGLE CYCLE NETWORK ANALYSIS

In the previous analysis using paired cycle routing (Section VI), our proposed generalization of R redundant quorums sets had moderate increases in resource usage and showed improvements to fault coverage. This section uses a similar experiment setup, however in contrast we are examining using additional quorum redundancy ($R = 3$) with just a single cycle compared to the pair of cycles used in prior art.

A. Fault-free operational analysis

The more links that a set of quorum cycles uses, the fewer (wavelength) resources that can be assigned to each link. Additionally each logical link represents a required physical transmitter and receiver, hence capital costs.

Table III shows significant 38.85 - 42.39% resource reduction when using $R = 3$ redundancy in quorums over the more traditional, prior art methods of simply using paired cycles. Using $R = 2$ gives even better resource reduction, as shown in Table III and [3]. This reduction represents the potential for lower capital costs in terms of physical transmitters and receivers needed and the potential for more (wavelength) resource availability within the network. The paired cycles results with a 95% confidence interval (CI) for $R = 1$ in Table I is repeated in column two of Table III for comparison to the single cycle, increased quorum redundancy technique. Our redundant quorum technique uses far fewer links (shown in columns 3 - 6.) $R = 2$ comes close to halving the necessary resources, whereas $R = 3$ is slightly larger at approximately 60% of the paired $R = 1$ solution.

Previously paired light-trails were used to form all of the point-to-point communication node pairs with minimum sized quorum cycles. This work and [3] consider utilizing intentionally formed redundant node pairs within the quorum routing to reduce the resources used as a potential trade-off to a small cost to network performance. We analyze the impact of increasing the redundancy within quorums to $R = 3$ and its impact of keeping resource utilization low. To measure this cost, metrics of missing node pairs is used.

Ideally, like the paired cycle case, there would be 0% missing, however single cycles don't have the benefit of both (a_i, a_j) and (a_j, a_i) pairs occurring in the same cycle. Table IV shows two important results. First, the dramatic reduction in resource utilization came at a trade off of a few missing communication pairs. $R = 2$ missed 0.90% or fewer on average (95% CI), and $R = 3$ missed even fewer at 0.26% or less on average (95% CI.) The paired cycles (column 2, Table III) used significantly more resources and did not miss any pairs (column 2, Table IV.) Secondly, compared to $R = 2$ single cycles, our redundant $R = 3$ cycles performs approximately 2+ times better every time. As seen in Table III, this performance improvement came at a only a slightly higher cost, while still being significantly smaller than the state of art approach.

The quorums set method guarantees that all of the pairs exist (Section V-B) It is the limitations of an unidirectional optical light-trails with its required one optical shutter in the off state per cycle that has caused the missing pairs and the potential need for additional compensation steps.

TABLE III. MEAN LINKS USED BY SINGLE CYCLES COMPARED TO PAIRED CYCLES USING OUR REDUNDANT QUORUM SOLUTION (95% CI.)

Networks	$R = 1$ (Paired)	$R = 2$ (Single)		$R = 3$ (Single)	
	Links	Links	Reduction	Links	Reduction
NSFNET	248.96 - 249.88	135.10 - 135.46	-45.76%	144.98 - 145.43	-41.79%
ARPANET	510.03 - 511.24	269.44 - 269.93	-47.19%	293.86 - 294.44	-42.39%
American	641.98 - 643.35	359.44 - 360.10	-44.02%	376.37 - 376.92	-41.39%
Chinese	2673.74 - 2678.28	1526.62 - 1528.85	-42.91%	1635.32 - 1637.26	-38.85%

TABLE IV. MEAN MISSING NODE PAIRS BY SINGLE CYCLES USING OUR REDUNDANT QUORUM SOLUTION (95% CI.)

Networks	$R = 1$ (Paired)	$R = 2$ (Single)	$R = 3$ (Single)
NSFNET	0.00%	0.81 - 0.90%	0.03 - 0.04%
ARPANET	0.00%	0.33 - 0.37%	0.10 - 0.13%
American	0.00%	0.50 - 0.54%	0.23 - 0.26%
Chinese	0.00%	0.26 - 0.27%	0.08 - 0.09%

Compensation is possible using an off-the-shelf solution of an additional routing step involving an Optical-to-Electrical-to-Optical (O/E/O) conversion and resending by a hub node. Even so, on average the $R = 2$ and $R = 3$ redundant quorums cycle solutions would require infrequent additional steps.

B. Fault-tolerant operational analysis

Using our generalized R quorum redundancy rather than cycle redundancy can save significant resources, however this cannot come at significant detriment to fault tolerance though.

1) *Single fault case:* Again to model the fault, we simulate the failure of each edge, $e \in E$, in the 1000 node mappings of each network model, $G = (V, E)$. We then examine the network's ability to serve all potential point-to-point requests by counting pairs of nodes that would be able to communicate and conversely those pairs that are unable to communicate. The results are then reported as fault coverage, total pairs able to communicate as a percentage of total point-to-point pairs. 100% would be perfect coverage, whereas 0% would be no fault coverage at all.

Our simulation results showed our redundant quorum-based cycle technique had 96.60 - 99.37% and 97.70 - 99.72% fault coverages, $R = 2$ and $R = 3$ respectively, in the four networks tested. In Figure 5, we compare the state-of-art paired cycle approach with our quorum redundant technique with single cycles that uses significantly fewer resources. With single edge failures, the paired cycles had a mean missing communication pair rate of less than 3 pairs or less than 0.53% across all networks (95% CI.) Hence the $R = 1$ (Paired) column shows mean fault coverage percentages is greater than 99.47% for all four networks. Our redundant quorum cycles technique, $R = 2$ and $R = 3$ (Single), could not reach that level of coverage, but did achieve an acceptable mean fault coverage rate (95% CI) of greater than 96.60 and 97.70% respectively across all networks.

While neither single cycle $R = 2$ or $R = 3$ could achieve the same level of fault coverage as the paired cycle solution, they did have missing rates better than 3.40 and 2.30% respectively, while achieving significant resource savings. In networks where an additional approximately 40% of resources could better be utilized for communication rather than redundancy, the trade off of missing a relatively small





Fig. 5. Percent mean fault coverage (95% CI) of our single cycle, redundant quorum solution experiencing a single link fault.

Fig. 6. Percent mean fault coverage (95% CI) of our single cycle, redundant quorum solution experiencing two simultaneous link faults. For graph clarity and consistency NSFNET's confidence intervals for $R = 2$ (Single) at 92.01 - 92.05% and for $R = 3$ (Single) at 93.23 - 93.27% mean fault coverage were not included in the graph.

percentage of communications during fault conditions may be considered tolerable.

2) *Two fault case:* The more complex two fault model considers all possible two edge failure combinations in the simulated networks. We then examine the network's ability to serve all potential point-to-point requests by counting pairs of nodes that would be able to communicate and conversely those pairs that are unable to communicate.

Our simulation results showed our redundant quorum-based cycle technique had 92.01 - 98.77% and 93.23 - 99.34% fault coverage, $R = 2$ and $R = 3$ respectively, in the four networks tested. In Figure 6, we compare the state-of-art paired cycle approach with our quorum redundant technique with single cycles that uses significantly fewer resources. With two edge failures, the paired cycles had a mean missing communication pair rate of less than 9 pairs or less than 2.37% across all networks (95% CI.) Hence in the $R = 1$ (Paired) column, it can be seen that the fault coverage is greater than 97.63%. Our redundant quorums set technique, $R = 2$ and $R = 3$ (Single), could not reach that level for all networks. Overall the missing pair rate (95% CI) was less than 7.99 and 6.77% respectively across all networks, which is reflected in fault coverages greater than 92.01%.

First, it is worth pointing out again that the paired cycles

solution uses more than 38% more resources on average (95% CI), while on average the redundant quorums cycle technique using only single cycles performed at most 6% worse in terms of mean fault coverage on the rarer two simultaneous fault case. This could be an acceptable trade off in many networks. Additionally, an interesting aspect of the data is that the results appear to have some network dependence. The networks with the larger number of nodes had better resiliency to faults than the smaller networks. This is likely a byproduct of our quorums set solution, where every cycle contains only a small, size K , subset of the total nodes. In all graphs, this translates into cycles that potentially can span the diameter of the graph. In larger graphs this could lead to more non-quorum nodes being passed through while forming the cycle, which could translate into more than the minimum calculated quorums set pairs (see Sections V-A and V-B.)

VIII. CONCLUSION

In this paper we proposed and evaluated a generalized formulation of redundant quorums sets for optical cycle routing. When the generalized R redundancy was applied to state of the art paired cycle techniques it provided optical networks almost fault-tolerant communications for both one and two simultaneous fault cases. When applied to our single cycle routing technique, the R redundant quorums significantly reduced resource usage and maintained high fault tolerance capabilities.

Quorums sets were chosen such that all network communication pairs appeared R times within a routed cycle set. We intentionally utilized this generalized R redundancy within the quorum cycles for fault-tolerance and reduction in resource usage. For $R = 2$ and $R = 3$, the paired cycle techniques on average used 5.63 - 14.18% and 15.21 - 22.29% more resources respectively, while on average achieving near fault-tolerance with 98.65 - 99.91% and 99.04 - 99.95% fault coverage rates respectively on the two simultaneous faults simulation. The single cycle technique had similar successes with almost fault-tolerant cycles that used significantly fewer resources (42.91 - 47.19% and 38.85 - 42.39% fewer respectively), while at the same time maintaining a high degree of fault-tolerance with 92.01 - 98.77% and 93.23 - 99.34% fault coverage respectively on the two simultaneous faults simulation.

In future work, we are examining ways to improve both the fault and fault-free performance further, while maintaining the significant resource usage reductions that the generalized R redundant quorums cycle solution provides.

ACKNOWLEDGMENT

Research funded in part by NSF Graduate Research Fellowship Program, IBM Ph.D. Fellowship Program, Symbi GK-12 Fellowship at Iowa State University, and the Virginia and Phillip Sproul Professorship at Iowa State University. The research reported in this paper is partially supported by the HPC@ISU equipment at Iowa State University, some of which has been purchased through funding provided by NSF under MRI grant number CNS 1229081 and CRI grant number 1205413. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the funding agencies.

REFERENCES

- [1] A. K. Somani and D. Lastine, "Optical paths supporting quorums for efficient communication," in *Optical Communications and Networks (ICOCN), 2014 13th International Conference on*. IEEE, 2014.
- [2] C. J. Kleinheksel and A. K. Somani, "Optical quorum cycles for efficient communication," *Photonic Network Communications*, pp. 1–10, 2015.
- [3] C. J. Kleinheksel and A. K. Somani, "Resource efficient redundancy using quorum-based cycle routing in optical networks," in *Transparent Optical Networks (ICTON), 2015 17th International Conference on*. IEEE, 2015, pp. 1–4.
- [4] I. Chlamtac and A. Gumaste, "Light-trails: A solution to ip centric communication in the optical domain," in *Quality of Service in Multiservice IP Networks*. Springer, 2003, pp. 634–644.
- [5] J. Fang, W. He, and A. K. Somani, "Optimal light trail design in wdm optical networks," in *Communications, 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 1699–1703.
- [6] W. D. Grover and G. Shen, "Extending the p-cycle concept to path-segment protection," in *Communications, 2003. ICC'03. IEEE International Conference on*, vol. 2. IEEE, 2003, pp. 1314–1319.
- [7] F. Zhang and W.-D. Zhong, "Performance evaluation of p-cycle based protection methods for provisioning of dynamic multicast sessions in mesh wdm networks," *Photonic Network Communications*, vol. 16, no. 2, pp. 127–138, 2008.
- [8] F. Zhang, W.-D. Zhong, and Y. Jin, "Optimizations of-cycle-based protection of optical multicast sessions," *Lightwave Technology, Journal of*, vol. 26, no. 19, pp. 3298–3306, 2008.
- [9] N. K. Singhal, L. H. Sahasrabudhe, and B. Mukherjee, "Provisioning of survivable multicast sessions against single link failures in optical wdm mesh networks," *Journal of lightwave technology*, vol. 21, no. 11, p. 2587, 2003.
- [10] Y. Qing and G. Ning, "Protecting dynamic multicast sessions in optical wdm mesh networks," in *Information, Communications and Signal Processing, 2005 Fifth International Conference on*. IEEE, 2005, pp. 1187–1190.
- [11] H. Luo, H. Yu, L. Li, and S. Wang, "On protecting dynamic multicast sessions in survivable mesh wdm networks," in *Communications, 2006. ICC'06. IEEE International Conference on*, vol. 2. IEEE, 2006, pp. 835–840.
- [12] F. Zhang and W.-D. Zhong, "Applying p-cycles in dynamic provisioning of survivable multicast sessions in optical wdm networks," in *Optical Fiber Communication Conference*. Optical Society of America, 2007, p. JWA74.
- [13] T. Feng, L. Ruan, and W. Zhang, "Intelligent p-cycle protection for multicast sessions in wdm networks," in *Communications, 2008. ICC'08. IEEE International Conference on*. IEEE, 2008, pp. 5165–5169.
- [14] A. Khalil, A. Hadjiantonis, G. Ellinas, and M. Ali, "Pre-planned multicast protection approaches in wdm mesh networks," in *Optical Communication, 2005. ECOC 2005. 31st European Conference on*. IET, 2005, pp. 25–26.
- [15] D. Lastine, S. Sankaran, and A. K. Somani, "A fault-tolerant multi-point cycle routing algorithm (mcra)," in *Broadband Communications, Networks, and Systems*. Springer, 2012, pp. 341–360.
- [16] A. Somani, D. Lastine, and S. Sankaran, "Finding complex cycles through a set of nodes," in *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, Dec 2011, pp. 1–5.
- [17] M. Maekawa, "An algorithm for mutual exclusion in decentralized systems," *ACM Transactions on Computer Systems (TOCS)*, vol. 3, no. 2, pp. 145–159, 1985.
- [18] W.-S. Luk and T.-T. Wong, "Two new quorum based algorithms for distributed mutual exclusion," in *Distributed Computing Systems, 1997., Proceedings of the 17th International Conference on*. IEEE, 1997, pp. 100–106.
- [19] L. Tang, W. Huang, M. Razo, A. Sivasankaran, M. Tacca, and A. Fumagalli, "Multicast tree computation in networks with multicast incapable nodes," in *High Performance Switching and Routing (HPSR), 2011 IEEE 12th International Conference on*. IEEE, 2011, pp. 95–100.